

Claims

[c1]

1. A method for in-order delivery of data within sequence from a first communication device to a second communication device in a system including a fabric, the method comprising:

receiving at the first communication device a data frame destined for the second communication device;

retrieving sequence information from said data frame, wherein data frames with the same sequence information require in-order delivery;

utilizing said sequence information to calculate a hash function; and

based on the calculated hash function, selecting one of a predetermined set of paths through the fabric connecting the first communication device to the second communication device.

[c2]

2. The method of claim 1, further comprising:

routing said data frame over the selected path to the second communication device.

[c3]

3. The method of claim 1, wherein the fabric is comprised of a plurality of interconnected Fibre Channel switches.

[c4]

4. The method of claim 1, wherein the first communication device is a Fibre Channel switch.

[c5]

5. The method of claim 1, wherein the second communication device is an end device in communication with the fabric.

[c6]

10059750-012502

6. The method of claim 1, wherein said sequence information includes information in at least one pre-selected field of a header of said data frame.

[c7]

7. The method of claim 6, wherein said hash function equals to the last digit of the sum of all values in said at least one pre-selected header field included in said sequence information.

[c8]

8. The method of claim 1, wherein said sequence information includes information in a source identifier field and a destination identifier field of a header of said data frame.

[c9]

9. The method of claim 1, wherein said sequence information includes information in at least one exchange identifier field of a header of said data frame.

[c10]

10. The method of claim 1, wherein each of said predetermined set of paths connecting the first communication device to the second communication device comprises a series of links between ports on adjacent communication devices.

[c11]

11. The method of claim 1, wherein each of said predetermined set of paths satisfies specified requirements on availability and on cost efficiency.

[c12]

12. The method of claim 1, wherein selecting one of a predetermined set of paths further comprises:

selecting one of a plurality of fields in an entry of a multiple-field routing table, wherein each field in the entry corresponds to one of said predetermined set of paths.

[c13]

13. The method of claim 12, wherein every value of the hash function is associated with a field in the entry of said multiple-field routing table.

[c14]

14. The method of claim 12, further comprising:

forwarding said data frame to a transmit port within the first communication device based on said selected field in the entry of said multiple-field routing table, wherein said selected path originates at said transmitting port.

10059760-012902

[c15]

15. A method for in-order delivery of data within sequence from a first communication device to a second communication device in a system including a fabric, the method comprising:

receiving at the first communication device a plurality of data frames destined for the second communication device;

retrieving sequence information from each of said plurality of data frames;

based on said retrieved sequence information, categorizing said plurality of data frames into a plurality of sequences, wherein each of said plurality of sequences requires in-order delivery; and

selecting for each of said plurality of sequences one of a predetermined set of paths through the fabric connecting the first communication device to the second communication device.

[c16]

16. The method of claim 15, further comprising:

routing said plurality of data frames over said predetermined set of paths, wherein all data frames belonging to a sequence use said selected path for said sequence.

[c17]

17. The method of claim 15, wherein said sequence information from each of said plurality of data frames includes information in at least one pre-selected field of a header of said data frame.

[c18]

18. The method of claim 15, wherein each of said plurality of sequences includes data frames with a same source device and a same destination device.

[c19]

19. The method of claim 15, wherein each of said plurality of sequences includes data frames within a same Fibre Channel exchange.

[c20]

20. The method of claim 15, wherein categorizing the plurality of data frames into a plurality of sequences further comprises:

associating an arbitrary number with all data frames corresponding to each of said plurality of sequences.

[c21]

21. The method of claim 20, wherein associating an arbitrary number with all data frames corresponding to each of the plurality of sequences further comprises:

associating an arbitrary number with each set of sequence information.

[c22]

22. The method of claim 21, wherein said arbitrary number corresponding to each set of sequence information is a hash function calculated from said set of sequence information.

[c23]

23. The method of claim 20, wherein selecting for each of the plurality of sequences one of a predetermined set of paths further comprises:

selecting, for each of said plurality of sequences, one of a predetermined set of paths based on said arbitrary number associated with the sequence.

[c24]

24. The method of claim 15, wherein selecting for each of the plurality of sequences one of a predetermined set of paths further comprises:

selecting, for each of said plurality of sequences, one of a plurality of fields in an entry to a multiple-field routing table, wherein each field in said entry corresponds to one of said predetermined set of paths.

[c25]

25. A method for in-order delivery of data within sequence from a first communication device to a second communication device in a system including a fabric, the method comprising:

identifying at least one header field as the basis for categorizing data frames into a plurality of sequences, wherein each sequence requires in-order delivery;

selecting a set of paths through the fabric connecting the first communication device to the second communication device; and

10059760-012902

constructing an entry of a multiple-field routing table, wherein each field in the entry corresponds to one of said selected set of paths and wherein at least one field is associated with each of said selected paths.

[c26]

26. The method of claim 25, further comprising:

upon receiving a data frame at the first communication device, routing the data frame to the second communication device based on header information of the data frame in said at least one header field, and the entry of said multiple-field routing table.

[c27]

27. The method of claim 26, wherein routing the data frame to the second communication device further comprises the steps performed at the first communication device of:

retrieving said header information from the data frame;

utilizing said header information to calculate a hash function;

choosing one of said selected paths based on said calculated hash function and said entry of said multiple-field routing table; and

routing the data frame over said selected path to the second communication device.

[c28]

28. The method of claim 25, wherein:

each of said plurality of sequences includes data frames with a same source device and a same destination device, and

said at least one header field includes a source identifier field and a destination identifier field.

[c29]

29. The method of claim 25, wherein:

each of said plurality of sequences includes data frames within a same Fibre Channel exchange, and

said at least one header field includes an originator exchange identifier field and a responder exchange identifier field.

[c30]

30. The method of claim 25, wherein selecting a set of paths through the fabric connecting the first communication device to the second communication device further comprises:

selecting a set of paths among all paths through the fabric connecting the first communication device to the second communication device based on specified requirements on availability and on cost efficiency.

[c31]

31. The method of claim 25, wherein each of a subset of said selected paths is associated with more than one field in the entry of said multiple-field routing table.

[c32]

32. The method of claim 31, wherein each path within the subset has higher bandwidth than any of the selected paths outside the subset.

[c33]

33. The method of claim 25, wherein each of a subset of said selected paths is associated with a weighting factor.

[c34]

34. A switch for in-order delivery of data within sequence through a fabric from a first communication device to a second communication device, the system comprising:

a data reception module for receiving from the first communication device a plurality of data frames destined for the second communication device;

a sequence identification module for retrieving sequence information from each of said plurality of data frames and for utilizing said retrieved sequence information to categorize said plurality of data frames into a plurality of sequences, wherein each of said plurality of sequences requires in-order delivery;

a path selection module for selecting for each of said plurality of sequences one of a predetermined set of paths through the fabric connecting the first communication device to the second communication device; and

a data transmission module for routing said plurality of data frames over said predetermined set of paths, wherein all data frames belonging to a sequence use said selected path for the sequence.

[c35]

10059760.012902

35. The switch of claim 34, wherein the sequence identification module further comprises:

a computation module for calculating a hash function from sequence information retrieved from each of said plurality of data frames; and

a data association module for associating said hash function calculated from said sequence information of each of said plurality of data frames with one of said plurality of sequences.

[c36]

36. The switch of claim 35, wherein the path selection module further comprises:

a path assignment module for assigning one of a predetermined set of paths to each of said plurality of sequences based on said calculated hash function associated with said sequence.

[c37]

37. The switch of claim 34, wherein the path selection module further comprises:

a multiple-field routing table including an entry for said selected set of paths, wherein each field in said entry corresponds to one of said selected set of paths and wherein at least one field is associated with each of said selected paths; and

a field assignment module for assignment one of said multiple fields of said entry to said multiple-field routing table to each of said plurality of sequences.

[c38]

38. The switch of claim 34, further comprising:

a preprocessing module for identifying at least one header field as the basis for the categorization of data frames into a plurality of sequences, and for selecting a set of paths through the fabric connecting the first communication device to the second communication device.

[c39]

39. A Fibre Channel network comprising:

a source device for providing a plurality of data frames comprising a sequence, each of said data frames including sequence information;

a target device for receiving said plurality of data frames from said source device;

a Fibre Channel fabric connecting said source and target devices, said Fibre Channel fabric including:

a first switch having an input coupled to said source device and having two outputs;

second and third switches, each of said second and third switches having an input coupled to one of said first switch outputs and having an output;

a fourth switch having two inputs, each coupled to one of the outputs of said second and third switches and having an output coupled to said target device, so that a sequence from said source device to said target device can be transmitted through either of said second or third switches,

wherein said first switch includes:

a data reception module for receiving from said source device a plurality of data frames destined for said target device;

a sequence identification module for retrieving sequence information from each of said plurality of data frames and for utilizing said retrieved sequence information to categorize said plurality of data frames into a plurality of sequences, wherein each of said plurality of sequences requires in-order delivery;

a path selection module for selecting for each of said plurality of sequences one of a predetermined set of paths through the fabric connecting said source device to said target device; and

a data transmission module for routing said plurality of data frames over said predetermined set of paths, wherein all data frames belonging to a sequence use said selected path for said sequence.

[c40]

40. The Fibre Channel network of claim 39, wherein the sequence identification module further comprises:

a computation module for calculating a hash function from sequence information retrieved from each of said plurality of data frames; and

a data association module for associating said hash function calculated from said sequence information of each of said plurality of data frames with one of said plurality of sequences.

[c41]

41. The Fibre Channel network of claim 40, wherein the path selection module further comprises:

a path assignment module for assigning one of a predetermined set of paths to each of said plurality of sequences based on said calculated hash function associated with said sequence.

[c42]

42. The Fibre Channel network of claim 39, wherein the path selection module further comprises:

a multiple-field routing table including an entry for said selected set of paths, wherein each field in said entry corresponds to one of said selected set of paths and wherein at least one field is associated with each of said selected paths; and

a field assignment module for assignment one of said multiple fields of said entry to said multiple-field routing table to each of said plurality of sequences.

[c43]

43. The Fibre Channel network of claim 39, further comprising:

a preprocessing module for identifying at least one header field as the basis for said categorization of data frames into a plurality of sequences, and for selecting a set of paths through the fabric connecting said source device to said target device.

10059760.012902